LCG and Tier-1 Facilities Status
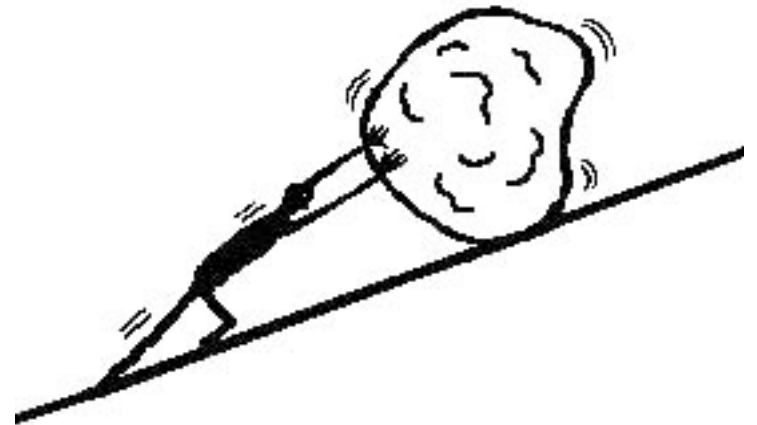
- LCG interoperability.

- Tier-1 facilities..

- Observations. (Not guaranteed to be wry, witty or nonobvious.)

© 2003 stephane chabrieres

LCG Interoperability

- Three interoperability choices

  1. Dedicated machinery .

     1. 5 nodes/10 cpus – now is 13 nodes/26 cpus no SE.
     2. LCG-2_0_0 – via LCFGng. Upgrade available.
     3. We have been "off the air" since #2 started.

  2. LCG gateway to US-CMS resources.

     1. LCG CE installed on cmssrv10 with condor-g – May 10.
     2. Sporadic work has been done:
        - problems – time ! = progress
     3. Solutions in new upgrade?

  3. Common interfaces.

Things to do

- Upgrade to current tag on dedicated resources and on CE interface (cmssrv10).

- Get back online with current resources.

  - Enter LCG fray.

- Configure CE interface.

  - Test simple jobs with current LCG.

  - Test simple jobs with one CMS cluster.

  - Configure farm to behave as LCG worker nodes. Test again.

SE needs to be installed

LCG Observations

- LCG support is good.

- LCG installation infrastructure is not (LCFGng).

- LCG instructions on doing manual installs is good. (Manual installs are Rocksible.)

- My input has been nil once up due to other time constraints.

- FNAL is a minimal presence due to time and resource constraints, both human and

Tier-1 Facilities – Current

- 196 machines in 5 clusters with various servers.

    - LCG: 13 WN's.

    - Storage: 20 dcache servers : admin, pool – 30 TB

    - 4 IBRIX nodes – 9 TB disk.

    - UAF – 3 servers, 36 WN's bunch o' disk.

    - PG – 4 servers, 84 WN's

    - DGT – 1 server, 4 WN's

    - Various servers – Rocks, console, bugzilla, VOMS, GIIS, GRIS, LCG-CE, MOP, tomcat/webserver, ganglia

Tier-1 Facilities – Future

- This year:
  - 100 WN's on the floor.
  - 4 new storage nodes, 5 new servers.
  - 10 TB storage.
- Next year
  - 300 new worker nodes up to 20 new servers.
  - Storage = gobs and gobs
- 3 FTE's

Stuff to do

- LCG interoperability.

- Dcache vs. Lustre vs. GFS

- Condor or some other batch system.

    - Sharing = fine grained policy

- Continued Rocks work – scalability, configuration etc.

- Load balancing for UAF – LVS, switch etc.

- High availability for critical servers.

- GRID3/Tier-2 involvement/support.

- Documentation.

- Monitoring – users, processes, machines.

- Security (FBS)

Observations

- Political priorities are trumped by physical realities.

- $\sum N(TE) <= \sum N(People)$

- CMS sysadmin ratio has been at 1:100 but many servers and storage with unique configurations included. (Google is at 1:500; Farms team is about 1:200-250 but they mostly deal with physical problems, no higher level issues, no research, not much grid - yet.)

- UAF, storage, PG are separate knowledge areas that need attention.